

# Binary Excess Risk for Smooth Convex Surrogates

Mehrdad Mahdavi

Michigan State University  
mahdavi@cse.msu.edu

Lijun Zhang

Michigan State University  
zhanglij@msu.edu

Rong Jin

Michigan State University  
rongjin@cse.msu.edu

## Abstract

In statistical learning theory, convex surrogates of the 0-1 loss are highly preferred because of the computational and theoretical virtues that convexity brings in. This is of more importance if we consider smooth surrogates as witnessed by the fact that the smoothness is further beneficial both computationally- by attaining an *optimal* convergence rate for optimization, and in a statistical sense- by providing an improved *optimistic* rate for generalization bound. In this paper we investigate the smoothness property from the viewpoint of statistical consistency and show how it affects the binary excess risk. We show that in contrast to optimization and generalization errors that favor the choice of smooth surrogate loss, the smoothness of loss function may degrade the binary excess risk. Motivated by this negative result, we provide a unified analysis that integrates optimization error, generalization bound, and the error in translating convex excess risk into a binary excess risk when examining the impact of smoothness on the binary excess risk. We show that under favorable conditions appropriate choice of smooth convex loss will result in a binary excess risk that is better than  $O(1/\sqrt{n})$ .

## 1 Introduction

In statistical binary classification, we are given an instance space  $\mathcal{X}$ , and a space  $\mathcal{Y} = \{-1, +1\}$  of labels representing which of two classes an instance belongs to. We assume that the product space  $\mathcal{X} \times \mathcal{Y}$  is endowed with a probability measure  $P_{\mathcal{X} \times \mathcal{Y}}$  that is unknown to the learner, however, it is possible to sample an arbitrary finite number of independent and identically distributed observations  $\mathcal{S} = ((\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)) \in (\mathcal{X} \times \mathcal{Y})^n$  from  $P_{\mathcal{X} \times \mathcal{Y}}$ . The goal is to output a classifier  $f : \mathcal{X} \mapsto \mathcal{Y}$  from a predefined class of functions  $\mathcal{F}$  that does well on unseen examples coming from the same distribution. The most natural loss function to measure the performance of a classifier is the misclassification error rate or the 0-1 (binary) loss  $\ell_{0-1}(f; (\mathbf{x}, y)) = \mathbb{I}[yf(\mathbf{x}) \leq 0]$ , where  $\mathbb{I}[z]$  is the indicator function that outputs 1 if  $z$  is true and zero, otherwise. We are interested, with the help of labeled examples  $\mathcal{S}$ , in finding a classifier  $f \in \mathcal{F}$  that minimizes the *binary excess risk*,

$$R(f) - R^* = \Pr_{(\mathbf{x}, y) \sim P_{\mathcal{X} \times \mathcal{Y}}} \{yf(\mathbf{x}) \leq 0\} - \min_g \Pr_{(\mathbf{x}, y) \sim P_{\mathcal{X} \times \mathcal{Y}}} \{yg(\mathbf{x}) \leq 0\}.$$

Many studies of binary excess risk assume that the optimal classifier  $f \in \mathcal{F}$  is learned by minimizing the empirical binary risk,  $\min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \mathbb{I}[y_i f(\mathbf{x}_i) \leq 0]$ , an approach that is usually referred to as Empirical Risk Minimization (ERM) [19]. To understand the generalization performance of the classifier learned by ERM, it is important to have upper bounds on the excess risk of the empirical minimizer that hold with a high probability and that take into account complexity measures of classification functions. It is well known that, under certain conditions, direct empirical classification error minimization is consistent [19] and achieves a fast convergence rate under low noise situations [14].

One shortcoming of the ERM based approaches is that they need to minimize 0-1 loss, leading to non-convex optimization problems that are potentially NP-hard<sup>1</sup> [1, 8]. A common practice to circumvent this difficulty is to replace the indicator function  $\mathbb{I}[\cdot \leq 0]$  with some convex loss  $\phi(\cdot)$  and find the optimal solution by minimizing the convex surrogate loss. Examples of such

<sup>1</sup>We note that several works [10, 11] provide efficient algorithms for direct 0-1 empirical error minimization but under strong (unrealistic) assumptions on data distribution or label generation.

surrogate loss functions for 0-1 loss include logit loss  $\phi_{\log}(f; (\mathbf{x}, y)) = \log(1 + \exp(-yf(\mathbf{x})))$  in logistic regression [7], hinge loss  $\phi_{\text{hing}}(f; (\mathbf{x}, y)) = \max(0, 1 - yf(\mathbf{x}))$  in support vector machine (SVM) [5] and exponential loss  $\phi_{\text{exp}}(f; (\mathbf{x}, y)) = \exp(-yf(\mathbf{x}))$  in AdaBoost [6].

It is known that under mild conditions, the classifier learned by minimizing the empirical loss of convex surrogate is consistent to the Bayes classifier [20, 13, 9, 12, 17, 2]. For instance, it was shown in [2] that the necessary and sufficient condition for a convex loss  $\phi(\cdot)$  to be consistent with the binary loss is that  $\phi(\cdot)$  is differentiable at origin and  $\phi'(0) < 0$ . It was further established in the same work that the binary excessive risk can be upper bound by the convex excess risk through a  $\psi$ -transform that depends on the surrogate convex loss  $\phi(\cdot)$ .

Since the choice of convex surrogates could significantly affect the binary excess risk, in this work, we will investigate the impact of the smoothness of a convex loss function on the binary excess risk. This is motivated by the recent results that show the advantages of using smooth convex surrogates in reducing the optimization complexity and the generalization error bound. More specifically, [15, 18] show that a faster convergence rate (i.e.  $O(1/T^2)$ ) can be achieved by first order methods when the objective function to be optimized is convex and smooth; in [16], the authors show that a smooth convex loss will lead to a better optimistic generalization error bound. Given the positive news of using smooth convex surrogates, an open research question is how the smoothness of a convex surrogate will affect the binary excess risk. The answer to this question, as will be revealed later, is negative: the smoother the convex loss, the poorer approximation will be for the binary excess risk. Thus, the second contribution of this work is to integrate these results for smooth convex losses, and examine the overall effect of replacing 0-1 loss with a smooth convex loss when taking into account three sources of errors, i.e. the optimization error, the generalization error, and the error in translating the convex excess risk into the binary risk. As we will show, under favorable conditions, appropriate choice of smooth convex loss will result a binary excess risk better than  $O(1/\sqrt{n})$ .

**Outline** The reminder of this paper is organized as follows. In Section 2 we set up notation and describe the setting. Section 3 briefly discusses the classification-calibrated convex surrogate losses on which our analysis relies. We derive the  $\psi$ -transform for smoothed hinge loss and elaborate its binary excess risk in Section 4. Section 5 provides a unified analysis of three types of errors and derives conditions in terms of smoothness to obtain better rates for the binary excess risk. Section 6 concludes the paper and the appendix contains the omitted proofs.

## 2 Problem Setting

Let  $\mathcal{S} = ((\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n))$  be a set of i.i.d. samples drawn from an unknown distribution  $P_{\mathcal{X} \times \mathcal{Y}}$  over  $\mathcal{X} \times \{-1, +1\}$ , where  $\mathbf{x}_i \in \mathcal{X} \subseteq \mathbb{R}^d$  is an instance and  $y_i \in \{-1, +1\}$  is the binary class assignment for  $\mathbf{x}_i$ . Let  $\kappa(\cdot, \cdot)$  be an universal kernel and let  $\mathcal{H}$  be the Reproducing Kernel Hilbert Space (RKHS) endowed with kernel  $\kappa(\cdot, \cdot)$ . According to [21],  $\mathcal{H}$  is a rich function space whose closure includes all the smooth functions. We consider predictors from  $\mathcal{H}$  with bounded norm to form the measurable function class  $\mathcal{F} = \{f \in \mathcal{H} : \|f\|_{\mathcal{H}} \leq B\}$ . Given a function  $f : \mathcal{X} \mapsto \mathbb{R}$ , define the risk of  $f$  as  $R(f) = \mathbb{E}_{(\mathbf{x}, y) \sim P_{\mathcal{X} \times \mathcal{Y}}} [\mathbb{I}[yf(\mathbf{x}) \leq 0]] = \Pr\{yf(\mathbf{x}) \leq 0\}$ . Let  $f^*$  be the optimal classifier that attains the minimum risk, i.e.  $f^* = \arg \min_f \Pr\{yf(\mathbf{x}) \leq 0\}$ . We assume  $f^* \in \mathcal{H}$  with  $\|f^*\|_{\mathcal{H}} \leq B$ . This boundedness condition is satisfied for any RKHS with a bounded kernel (i.e.  $\sup_{\mathbf{x} \in \mathcal{X}} \kappa(\mathbf{x}, \mathbf{x}) \leq B$ ). Henceforth, let  $R^*$  stand for the minimum achievable risk by the optimal classifier  $f^*$ , i.e.,  $R^* = R(f^*)$ . Define the *binary excess risk* for a prediction function  $f \in \mathcal{F}$  as

$$\mathcal{E}(f) = R(f) - R^*.$$

Our goal is to efficiently learn a prediction function  $f \in \mathcal{F}$  from the training examples in  $\mathcal{S}$  that minimizes the binary excess risk  $\mathcal{E}(f)$ . As already mentioned, for computational convenience, we work with a convex loss function rather than the 0-1 loss function. Given a convex surrogate loss function  $\phi : \mathbb{R} \mapsto \mathbb{R}_+$  (e.g., hinge loss, exponential loss, or logistic loss) we define the risk with respect to the convex loss  $\phi$  (convex risk or  $\phi$ -risk) as

$$R_{\phi}(f) = \mathbb{E}_{(\mathbf{x}, y) \sim P_{\mathcal{X} \times \mathcal{Y}}} [\phi(yf(\mathbf{x}))].$$

Similarly we define the *optimal*  $\phi$ -risk as  $R_\phi^* = \inf_{f \in \mathcal{F}} \mathbb{E}_{(\mathbf{x}, y) \sim P_{\mathcal{X} \times \mathcal{Y}}}[\phi(yf(\mathbf{x}))]$ . The *excess  $\phi$ -risk* or *convex excess risk* of a classifier  $f \in \mathcal{F}$  with respect to the convex surrogate loss  $\phi(\cdot)$  is defined as

$$\mathcal{E}_\phi(f) = R_\phi(f) - R_\phi^*.$$

An important line of research in statistical learning theory focused on relating the convex excess risk  $\mathcal{E}_\phi(f)$  to the binary excess risk  $\mathcal{E}(f)$  that will be elaborated in next section.

### 3 Classification Calibration and Surrogate Risk Bounds

Although it is computationally convenient to minimize the empirical risk based on a convex surrogate, the ultimate goal of any classification method is to find a function  $f$  that minimizes the binary loss. Therefore, it is crucial to investigate the conditions which guarantee that if the  $\phi$ -risk of  $f$  gets close to the optimal  $R_\phi^*$ , the binary risk of  $f$  will also approach the optimal binary risk  $R^*$ . This question has been an active trend in statistical learning theory over the last decade where the necessary and sufficient conditions have been established for relating the binary excess risk to a convex excess risk [20, 13, 9, 12, 17, 2].

In this paper we follow the strategy introduced in [2] in order to relate the binary excess risk to the excess  $\phi$ -risk. Their methodology, through the notion of *classification calibration*, allows us to find quantitative relationship between the excess risk associated with  $\phi$  and the excess risk associated with 0-1 loss. It is established in [2] that the binary excessive risk can be bounded by the convex excess risk, based on the convex loss function  $\phi$ , through a  $\psi$ -transform.

**Definition 1.** Given a loss function  $\phi : \mathbb{R} \mapsto [0, \infty)$ , define the function  $\psi : [0, 1] \mapsto [0, \infty)$  by

$$\tilde{\psi}(z) = H^- \left( \frac{1+z}{2} \right) - H \left( \frac{1+z}{2} \right)$$

where

$$H^-(\eta) = \inf_{\alpha: \alpha(2\eta-1) \leq 0} (\eta\phi(\alpha) + (1-\eta)\phi(-\alpha)) \text{ and } H(\eta) = \inf_{\alpha \in \mathbb{R}} (\eta\phi(\alpha) + (1-\eta)\phi(-\alpha)).$$

The transform function  $\psi : [0, 1] \mapsto [0, \infty)$  is defined to be the convex closure of  $\tilde{\psi}$ .

The following theorem from [2, Theorem 1] shows that the binary excess risk can be bounded by the convex excess risk using transform function  $\psi : [0, 1] \mapsto [0, \infty)$  that depends on the surrogate convex loss function.

**Theorem 1.** For any non-negative loss function  $\phi(\cdot)$ , any measurable function  $f \in \mathcal{F}$ , and any probability distribution  $P_{\mathcal{X} \times \mathcal{Y}}$  on  $\mathcal{X} \times \mathcal{Y}$ , there is a nondecreasing function  $\psi : [0, 1] \mapsto [0, \infty)$  that

$$\psi(R(f) - R^*) \leq R_\phi(f) - R_\phi^* \quad (1)$$

holds. Here the minimization is taken over all measurable functions.

**Definition 2.** A convex loss  $\phi$  is *classification-calibrated* if, for any  $\eta \neq 1/2$ ,

$$H^-(\eta) > H(\eta).$$

This condition is essentially an extension of [20, Theorem 2.1] and can be viewed as a form of Fisher consistency that is appropriate for classification.

It has been shown in [2] that the necessary and sufficient condition for a convex loss  $\phi(z)$  to be classification-calibrated is if it is differentiable at the origin and  $\phi'(0) < 0$ . In particular, for a certain convex function  $\phi(\cdot)$ , the  $\psi$ -transform can be computed by

$$\psi(z) = \inf_{\alpha \leq 0} \left( \frac{1+z}{2} \phi(\alpha) + \frac{1-z}{2} \phi(-\alpha) \right) - \inf_{\alpha \in \mathbb{R}} \left( \frac{1+z}{2} \phi(\alpha) + \frac{1-z}{2} \phi(-\alpha) \right),$$

that can be further simplified as  $\psi(z) = \phi(0) - H\left(\frac{1+z}{2}\right)$  when  $\phi$  is classification-calibrated. Examples of  $\psi$ -transform for the convex surrogate functions of known practical algorithms mentioned before are as follows: (i) for hinge loss  $\phi(\alpha) = \max(0, 1 - \alpha)$ ,  $\psi(z) = |z|$ , (ii) for exponential loss  $\phi(\alpha) = e^{-\alpha}$ ,  $\psi(z) = 1 - \sqrt{1 - z^2} \geq z^2/2$ , and (iii) for truncated quadratic loss  $\phi(\alpha) = [\max(0, 1 - \alpha)]^2$ ,  $\phi(z) = z^2$ .

**Remark 1.** We note that the inequality in (1) provides insufficient guidance on choosing appropriate loss function. First, it does not measure explicitly how the choice of the convex surrogate  $\phi(\cdot)$  affects the excess risk  $R_\phi(f) - R_\phi^*$ . Second, it does not take into account the impact of loss function on optimization efficiency, an important issue for practitioners when dealing with big data. It is thus unclear, from Theorem 1, how to choose an appropriate loss function that could result in a small generalization error for the binary loss when the computational time is limited. In this paper, we address these limitations by examining a family of convex losses that are constructed by smoothing the hinge loss function using different smoothing parameters. We study the binary excessive risk of the learned classification function by taking into account errors in optimization, generalization, and translation of convex excess risk into binary excess risk.

## 4 Binary Excess Risk for Smoothed Hinge Loss

As stated before, to efficiently learn a prediction function  $f \in \mathcal{F}$ , we will replace the binary loss with a smooth convex loss. Since hinge loss is one of the most popular loss functions used in machine learning and is the loss of choice for classification problems in terms of the margin error [3], in this work, we will focus on the smoothed version of the hinge loss. Another advantage of using the hinge loss is that its  $\psi$ -transform is a linear function. Compared with the  $\psi$ -transforms of other popular convex loss functions (e.g. exponential loss and truncated square loss) that are mostly quadratic, using the hinge loss as convex surrogate will lead to a tighter bound for the binary excess risk.

The smoothed hinge loss considered in this paper is defined as

$$\phi(z; \gamma) = \max_{\alpha \in [0,1]} \alpha(1-z) + \frac{1}{\gamma} \mathcal{R}(\alpha), \quad (2)$$

where  $\mathcal{R}(\alpha) = -\alpha \log \alpha - (1-\alpha) \log(1-\alpha)$  and  $\gamma > 0$  is the smoothing parameter. It is straightforward to verify that the loss function in (2) can be simplified as

$$\phi(z; \gamma) = \frac{1}{\gamma} \log(1 + \exp(\gamma(1-z))).$$

It is not immediately clear from Theorem 1 how the relationship between smooth convex excess risk  $\mathcal{E}_\phi(\cdot)$  and binary excess risk is affected by the smoothness parameter  $\gamma$ . In addition, as discussed in [2], whereas conditions such as convexity and smoothness have natural relationship to optimization and generalization, it is not immediately obvious how properties such as convexity and smoothness of convex surrogate relates to statistical consequences. In what follows, we show that, indeed smoothness of loss function has a *negative* statistical consequence and can degrade the binary excess risk.

### 4.1 $\psi$ -Transform for smoothed hinge loss

The first step in our analysis is to derive the  $\psi$ -transform for the loss function defined in (2) as stated in the following theorem.

**Theorem 2.** The  $\psi$ -transform of smoothed hinge loss with smoothing parameter  $\gamma$  is given by

$$\psi(\eta; \gamma) = -\frac{1+\eta}{2\gamma} \log \left( \frac{1}{1+e^\gamma} \left[ 1 + e^\gamma \frac{C_1}{1+\eta} \right] \right) - \frac{1-\eta}{2\gamma} \log \left( \frac{1}{1+e^\gamma} \left[ 1 + e^\gamma \frac{C_2}{1-\eta} \right] \right)$$

where  $C_1$  and  $C_2$  are defined as  $C_1 = -\eta e^\gamma + \sqrt{\eta^2 e^{2\gamma} + 1 - \eta^2}$  and  $C_2 = \eta e^\gamma + \sqrt{\eta^2 e^{2\gamma} + 1 - \eta^2}$ .

The  $\psi$ -transform given in Theorem 2 is too complicated to be useful. The theorem below provides a simpler bound for the  $\psi$ -transform in terms of the smoothness parameter  $\gamma$ .

**Theorem 3.** For  $\eta \in (-1, 1)$ , we have

$$\psi(\eta; \gamma) \geq |\eta| - \frac{1}{\gamma} \log \frac{1}{|\eta|}.$$

**Remark 2.** The bound obtained in Theorem 3 demonstrates that when  $\gamma$  approaches to infinity, the  $\psi$ -transform for smoothed hinge loss  $\phi(\eta; \gamma)$  becomes  $|\eta|$ . According to [2], the  $\psi$ -transform for the hinge loss is  $\psi(\eta) = |\eta|$ . Therefore, this result is consistent with the  $\psi$ -transform for smoothed hinge loss, which is the limit of  $\phi(z; \gamma)$  as  $\gamma$  approaches infinity.

## 4.2 A bound on binary excess risk $\mathcal{E}(f)$ based on smooth convex excess risk $\mathcal{E}_\phi(f)$

Based on the transform function  $\psi(\cdot; \gamma)$  that is computed for smoothed hinge loss with smoothing parameter  $\gamma$ , we are now in the position to bound its corresponding binary excess risk  $\mathcal{E}(f)$ . Our main result in this section is the following theorem that shows how binary excess risk can be bounded by the excess  $\phi$ -risk for smoothed hinge loss.

**Theorem 4.** Consider any measurable function  $f \in \mathcal{F}$  and the smoothed hinge loss  $\phi(\cdot)$  with parameter  $\gamma$  defined in (2). Then, binary excess risk  $\mathcal{E}(f)$  can be bounded by the smooth convex excess risk  $\mathcal{E}_\phi(f)$  as

$$\mathcal{E}(f) \leq \mathcal{E}_\phi(f) + \frac{\mathcal{E}_\phi(f)}{1 + \gamma \mathcal{E}_\phi(f)} \log \frac{1}{\mathcal{E}_\phi(f)}.$$

*Proof.* Using the result from Theorem 1, we have  $\mathcal{E}_\phi(f) \geq \psi(\mathcal{E}(f); \gamma)$  and therefore an immediate result from the  $\psi$ -transform for smoothed hinge loss that is obtained in Theorem 3 indicates

$$\mathcal{E}(f) + \frac{1}{\gamma} \log \mathcal{E}(f) \leq \mathcal{E}_\phi(f).$$

Define  $\Delta = \mathcal{E}(f) - \mathcal{E}_\phi(f)$ . We have

$$\Delta + \frac{1}{\gamma} \log(\Delta + \mathcal{E}_\phi(f)) = \Delta + \frac{1}{\gamma} \log \mathcal{E}_\phi(f) + \frac{1}{\gamma} \log \left( 1 + \frac{\Delta}{\mathcal{E}_\phi(f)} \right) \leq 0.$$

Based on the  $\log(1+x) \leq x$  inequality, the sufficient condition for the above inequality to hold is to have

$$\Delta + \frac{\Delta}{\gamma \mathcal{E}_\phi(f)} \leq \frac{1}{\gamma} \log \frac{1}{\mathcal{E}_\phi(f)}$$

and therefore

$$\Delta \leq \frac{\gamma^{-1}}{1 + (\gamma \mathcal{E}_\phi(f))^{-1}} \log \frac{1}{\mathcal{E}_\phi(f)} = \frac{\mathcal{E}_\phi(f)}{1 + \gamma \mathcal{E}_\phi(f)} \log \frac{1}{\mathcal{E}_\phi(f)}.$$

The final bound is obtained by substituting  $\mathcal{E}(f) - \mathcal{E}_\phi(f)$  for  $\Delta$  in the left hand side of above inequality.  $\square$

As indicated by Theorem 4, the smaller the smoothing parameter  $\gamma$ , the poorer the approximation is in bounding the binary excess  $\mathcal{E}(f)$  with smooth convex excess risk  $\mathcal{E}_\phi(f)$ . On the other hand, the smoothness of loss function has been proven to be beneficial in terms of optimization error and generalization bound. The mixture of negative and positive results for using smooth convex surrogates motivates us to develop an integrated bound for binary excess risk that takes into account all types of errors. One of the main contributions of this work is to show that under favorable conditions, with appropriate choice of smoothing parameter, the smoothed hinge loss will result in a bound for the binary excess risk better than  $O(1/\sqrt{n})$ .

## 5 A Unified Analysis of Binary Excess Risk for Smooth Surrogates

Using the smoothed hinge loss, we define the convex loss for a prediction function  $f \in \mathcal{F}$  as  $R_\phi(f) = \mathbb{E}[\phi(yf(\mathbf{x}); \gamma)]$ . Let  $f_\gamma^*$  be the optimal classifier that minimizes  $R_\phi(f)$ . Similar to the case of binary loss, we assume  $f_\gamma^* \in \mathcal{H}$  with  $\|f_\gamma^*\| \leq B$ . The smooth convex excess risk for a given prediction function  $f \in \mathcal{F}$  is then given by  $\mathcal{E}_\phi(f) = R_\phi(f) - R_\phi(f_\gamma^*)$ . Given the smooth

convex loss  $\phi(z; \gamma)$  in (2), we find the optimal classifier by minimizing the empirical convex loss, i.e.  $\min_{f \in \mathcal{H}, \|f\|_{\mathcal{H}} \leq B} \widehat{R}_\phi(f)$ , where the empirical convex loss  $\widehat{R}_\phi(f)$  is given by

$$\widehat{R}_\phi(f) = \frac{1}{n} \sum_{i=1}^n \phi(y_i f(\mathbf{x}_i); \gamma). \quad (3)$$

Let  $\widehat{f}$  be the solution learned from solving the empirical convex loss over training examples. There are three sources of errors that affect bounding the binary excess risk  $\mathcal{E}(\widehat{f})$ . First, since  $\widehat{f}$  is obtained by numerically solving an optimization problem, the error in estimating the optimal solution, which we refer to as optimization error<sup>2</sup>, will affect  $\mathcal{E}(\widehat{f})$ . Additionally, since the binary excess risk can be bounded by a nonlinear transform of the convex excess risk, both the bound for  $\mathcal{E}_\phi(\widehat{f})$  and the error in approximating  $\mathcal{E}(\widehat{f})$  with  $\mathcal{E}_\phi(\widehat{f})$  will affect the final estimation of  $\mathcal{E}(\widehat{f})$ . We aim at investigating how the smoothing parameter  $\gamma$  affect all these three types of errors. As it is investigated in Theorem 4, a smaller smoothing parameter  $\gamma$  will result in a poorer approximation of  $\mathcal{E}(\widehat{f})$ . On the other hand, a smaller smoothing parameter  $\gamma$  will result in a smaller estimation error and a smaller bound for  $\mathcal{E}_\phi(\widehat{f})$ . Based on the understanding of how smoothing parameter  $\gamma$  affects the three errors, we identify the choice of  $\gamma$  that results in the best tradeoff between all three error and consequentially a binary excess risk  $\mathcal{E}(\widehat{f})$  better than  $O(1/\sqrt{n})$ .

To investigate how the smoothing parameter  $\gamma$  affects the binary excess risk  $\mathcal{E}(\widehat{f})$ , we intend to unify three types of errors. The analysis is comprised of two components, i.e. bounding the binary excess risk  $\mathcal{E}(f)$  by a smooth convex excess risk  $\mathcal{E}_\phi(f)$  that has been established in Theorem 4 and bounding  $\mathcal{E}_\phi(f)$  for a solution  $f$  that is suboptimal in minimizing the empirical convex loss  $\widehat{R}_\phi(f)$  that is the focus of this section.

### 5.1 Bounding smooth excess convex risk $\mathcal{E}_\phi(f)$

We now turn to bounding the excess  $\phi$ -risk  $\mathcal{E}_\phi(f)$  for the smoothed hinge loss. To bound  $\mathcal{E}_\phi(f)$  we need to consider two types of errors: optimization error due to the approximate optimization of the empirical  $\phi$ -risk, and the generalization error bound for the empirical risk minimizer. After obtaining these two errors for smooth convex surrogates, we provide a unified bound on the excess  $\phi$ -risk  $\mathcal{E}_\phi(f)$  of empirical convex risk minimizer in terms of  $n$ .

We begin by bounding the error arising from solving the optimization problem numerically. One nice property of smoothed hinge loss function is that both its first order and second order derivatives are bounded, i.e.

$$|\phi'(z; \gamma)| = \left| \frac{\exp(\gamma(1-z))}{1 + \exp(\gamma(1-z))} \right| \leq 1, \quad \phi''(z; \gamma) = \gamma \frac{\exp(\gamma(1-z))}{(1 + \exp(\gamma(1-z)))^2} \leq \frac{\gamma}{4}.$$

Due to the smoothness of  $\phi(z; \gamma)$ , we can apply the accelerated optimization algorithm [15, 18] to achieve an  $O(1/k^2)$  convergence rate for the optimization, where  $k$  is the number of iterations the optimization algorithm proceeds. More specifically, we will apply Algorithm 1 from [18] to solve the numerical optimization problem in (3) over the convex domain  $\mathcal{F} = \{f \in \mathcal{H} : \|f\|_{\mathcal{H}} \leq B\}$  which results in the following updating rules at  $s$ th iteration:

$$\begin{aligned} g_s &= (1 - \theta_s)f_s + \theta_s h_s \\ h_{s+1} &= \arg \min_{h \in \mathcal{F}} \left( \langle \nabla \widehat{R}_\phi(g_s), h - g_s \rangle + \frac{\theta_s}{2} \|h - h_s\|_{\mathcal{H}} \right) \\ f_{s+1} &= (1 - \theta_s)f_s + \theta_s h_{s+1}. \end{aligned} \quad (4)$$

The following theorem that follows immediately from [18, Corollary 1] and the fact  $\phi''(z; \gamma) \leq \gamma/4$ , bounds the optimization error for the optimization problem after  $k$  iterations.

<sup>2</sup>We note that in literature the error in estimating the optimal solution for empirical minimization is usually referred to as *estimation error*. We emphasize it as optimization error because different convex surrogates could lead to very different iteration complexities and consequentially different optimization efficiency.



**Lemma 1.** Let  $\hat{f} = f_{k+1}$  be the solution obtained by running accelerated gradient descent method (i.e., updating rules in (4)) to solve the optimization problem in (3) after  $k$  iterations with  $\theta_0 = 1$  and  $\theta_k = 2/(k+2)$  for  $k \geq 1$ . We have

$$\hat{R}_\phi(\hat{f}) \leq \min_{\|f\|_{\mathcal{H}} \leq B} \hat{R}_\phi(f) + \frac{\gamma B^2}{(k+2)^2}.$$

We now turn to understanding the generalization error for the smooth convex loss. There are many theoretical results giving upper bounds of the generalization error. However, a recent result [16] has showed that it is possible to obtain optimistic rates for generalization bound of smooth convex loss (in the sense that smooth losses yield better generalization bounds when the problem is easier), which are more appealing than the generalization of simple Lipschitz continuous losses. The following theorem from [16, Theorem 1] bounds the generalization error for any solution  $f \in \mathcal{F}$  when the learning has been performed by a smooth convex surrogate  $\phi(\cdot)$ .

**Lemma 2.** With a probability  $1 - \delta$ , for any  $\|f\|_{\mathcal{H}} \leq B$ , we have

$$\begin{aligned} R_\phi(f) - \hat{R}_\phi(f) &\leq K_1 \left( \frac{(B + \gamma B^2)t}{n} + \sqrt{\hat{R}_\phi(f) \frac{(B + \gamma B^2)t}{n}} \right) \\ R_\phi(f) - \hat{R}_\phi(f) &\leq K_2 \left( \frac{(B + \gamma B^2)t}{n} + \sqrt{R_\phi(f) \frac{(B + \gamma B^2)t}{n}} \right). \end{aligned}$$

where  $t = \log(1/\delta) + \log^3 n$  and  $K_1$  and  $K_2$  are universal constants.

The bound stated in this lemma is *optimistic* in the sense that it reduces to  $\tilde{O}(1/\sqrt{n})$  when the problem is difficult and be better when the problem is easier, approaching  $\tilde{O}(1/n)$  for linearly separable data, i.e.,  $R_\phi^* = 0$  in the second inequality. These two lemmas essentially enable us to transform a bound on the optimization error and generalization bound into a bound on the convex excess risk. In particular, by combining Lemma 1 with Lemma 2, we have the following theorem that bounds the smooth convex excess risk  $\mathcal{E}_\phi(\hat{f}) = R_\phi(\hat{f}) - R_\phi(f_\lambda^*)$  for the empirical convex risk minimizer.

**Theorem 5.** Let  $\hat{f}$  be the solution output from updating rules in (4) after  $k$  iterations. Then, with a probability at least  $1 - \delta$ , we have

$$\mathcal{E}_\phi(\hat{f}) \leq \frac{\gamma B^2}{(k+2)^2} + K \left( \frac{(B + \gamma B^2)t}{n} + \sqrt{R_\phi^* \frac{(B + \gamma B^2)t}{n}} + \sqrt{\frac{\gamma B^2(B + \gamma B^2)t}{(k+2)^2 n}} \right)$$

where  $K$  is a universal constant,  $t = \log(1/\delta) + \log^3 n$ , and  $R_\phi^* = \min_{\|f\|_{\mathcal{H}} \leq B} R_\phi(f)$ .

Since our overall interest is to understand how the smoothing parameter  $\gamma$  affects the convergence rate of excess risk in terms of  $n$ , the number of training examples, it is better to parametrize both the number of iterations  $k$  and smoothing parameter  $\gamma$  in  $n$ , and bound the  $\mathcal{E}_\phi(\hat{f})$  only in terms of  $n$ . This is given in the following corollary.

**Corollary 6.** Assume  $\gamma \geq 1$  and  $B \geq 1$ . Parametrize  $k$  and  $\gamma$  in terms of  $n$  as  $k+2 = n^{\alpha/2}$  and  $\gamma = n^\beta$ . Then, with a probability at least  $1 - \delta$ ,

$$\mathcal{E}_\phi(\hat{f}) \leq C(B, t) \left( n^{\beta-\alpha} + n^{\beta-1} + n^{\beta-(\alpha+1)/2} + [R_\phi^*]^{1/2} n^{(\beta-1)/2} \right) \quad (5)$$

where  $C(B, t)$  is a constant depending on both  $B$  and  $t$  with  $t = \log(1/\delta) + \log^3 n$ .

The bound given in (5) depends on  $R_\phi^*$ . We would like to further characterize  $R_\phi^*$  in terms of  $\gamma$ . First, we have

$$\begin{aligned} \phi(z; \gamma) &= \max_{\alpha \in [0, 1]} \max(0, 1 - z) + \frac{1}{\gamma} \mathcal{R}(\alpha) \\ &\leq \max_{\alpha \in [0, 1]} \max(0, 1 - z) + \frac{1}{\gamma} \log 2 = \phi_{\text{hinge}}(z) + \frac{\log 2}{\gamma}, \end{aligned}$$

where  $\phi_{\text{hinge}}(z) = \max(0, 1 - z)$  is the hinge loss. As a result, we have

$$R_{\phi}^* \leq R_{\text{hinge}}^* + \frac{\log 2}{\gamma}$$

where  $R_{\text{hinge}}^* = \min_{\|f\|_{\mathcal{H}} \leq B} \mathbb{E}_{(\mathbf{x}, y) \sim P_{\mathcal{X} \times \mathcal{Y}}} [\phi_{\text{hinge}}(yf(\mathbf{x}))]$  is the optimal risk with respect to the hinge loss. In general, we will assume

$$R_{\phi}^* \leq R_{\text{hinge}}^* + \frac{a}{\gamma^{1+\xi}} \quad (6)$$

where  $a > 0$  is a constant and  $\xi \geq 0$  characterizes how fast  $R_{\phi}^*$  will converge to  $R_{\text{hinge}}^*$  with increasing  $\gamma$ . To see why the assumption in (6) is sensible, consider the case when the optimal classifier  $f_{\text{hinge}}^* = \arg \min_{\|f\|_{\mathcal{H}} \leq B} R_{\text{hinge}}(f)$  can perfectly classify all the data points with margin  $\epsilon$ , in which we have

$$R_{\phi}^* \leq R_{\text{hinge}}^* + O\left(\frac{e^{-\epsilon\gamma}}{\gamma}\right)$$

which satisfy the condition in (6) with arbitrarily large  $\xi$ . It is easy to verify that the condition (6) holds with  $\xi > 0$  if  $f_{\text{hinge}}^*$  can perfectly classify  $O(1 - \gamma^{-1-\xi})$  percentage of data with margin  $\epsilon$ .

Using the assumption in (6), we have the following result that characterizes the smooth convex excess risk bound  $\mathcal{E}_{\phi}(\hat{f})$  stated in terms of the parameters  $\alpha$ ,  $\delta$  and  $R_{\text{hinge}}^*$ .

**Theorem 7.** Assume  $\alpha \geq 1/2$ . Set  $\beta$  as

$$\beta = \frac{\min(1/2, \alpha - 1/2)}{1 + \xi}.$$

With a probability  $1 - \delta$ , we have

$$\mathcal{E}_{\phi}(\hat{f}) \leq O(n^{-\tau_1} + [R_{\text{hinge}}^*]^{1/2} n^{-\tau_2})$$

where

$$\tau_1 = \frac{1 + 2\xi \min(1, \alpha)}{2(1 + \xi)}, \quad \tau_2 = \frac{1/2 + \xi}{2(1 + \xi)}$$

*Proof.* Replacing  $R_{\phi}^*$  in Corollary 6 with the expression in (6), we have, with a probability  $1 - \delta$ ,

$$\mathcal{E}_{\phi}(\hat{f}) \leq C(R, t, a) \left( n^{\beta-\alpha} + n^{\beta-1} + n^{\beta-(\alpha+1)/2} + [R_{\text{hinge}}^*]^{1/2} n^{(\beta-1)/2} + n^{-1/2-\xi\beta} \right)$$

We first consider the case when  $\alpha > 1$ . In this case, we have

$$\mathcal{E}_{\phi}(\hat{f}) \leq O\left(n^{\beta-1} + n^{-1/2-\xi\beta} + [R_{\text{hinge}}^*]^{1/2} n^{(\beta-1)/2}\right)$$

By choosing  $\beta - 1 = -1/2 - \xi\beta$ , we have  $\beta = \frac{1/2}{1+\xi}$  and

$$\mathcal{E}_{\phi}(\hat{f}) \leq O(n^{-(1/2+\xi)/(1+\xi)} + [R_{\text{hinge}}^*]^{1/2} n^{-(1/2+\xi)/[2(1+\xi)]})$$

In the second case, we have  $\alpha \in [1/2, 1]$ . Hence we have

$$\mathcal{E}_{\phi}(\hat{f}) \leq O\left(n^{\beta-\alpha} + [R_{\text{hinge}}^*]^{1/2} n^{(\beta-1)/2} + n^{-1/2-\xi\beta}\right)$$

By setting  $\beta - \alpha = -1/2 - \xi\beta$ , we have  $\beta = \frac{\alpha-1/2}{1+\xi}$  and

$$\mathcal{E}_{\phi}(\hat{f}) \leq O\left(n^{-\frac{\xi\alpha+1/2}{1+\xi}} + [R_{\text{hinge}}^*]^{1/2} n^{-(1/2+\xi)/[2(1+\xi)]}\right).$$

We complete the proof by combining the results for the two cases.  $\square$



## 5.2 Bounding binary excess risk $\mathcal{E}(f)$

We now combine the results from Theorem 4 and Corollary 6 to bound  $\mathcal{E}(f)$ .

**Theorem 8.** Assume  $\alpha \geq 1/2$ . For a failure probability  $\delta \in (0, 1)$ , define  $n_0$  as

$$n_0 \leq K_3(B, \delta) \left( \frac{1}{R_{\text{hinge}}^*} \right)^{1/(2\tau_1 - 2\tau_2)}$$

where  $K_3(B, \delta)$  is a constant depending on  $B$  and  $\delta$ , and  $\tau_1$  and  $\tau_2$  are defined in Theorem 7. Set  $\beta$  as that in Theorem 7 if  $n \leq n_0$  and 0, otherwise. Then, with a probability  $1 - \delta$ , we have

$$\mathcal{E}_\phi \leq \begin{cases} K_4(B, \delta) n^{-\tau_1} \log n & n \leq n_0 \\ K_5(B, \delta) n^{-1/2} \log n & n > n_0 \end{cases}$$

where  $K_4(B, \delta)$  and  $K_5(B, \delta)$  are constants depending on  $B$  and  $\delta$ .

Theorem 8 follows from Theorem 4 and similar analysis for Theorem 7, from which we have

$$\mathcal{E}(\hat{f}) = R(\hat{f}) - R_* \leq O \left( \min \left( \gamma^{-1}, R_\phi(\hat{f}) - R_\phi^* \right) \log n \right)$$

**Remark 3.** According to Theorem 8, when the number of training examples  $n$  is not too large, for the binary excess risk of empirical minimizer we have, with a high probability,

$$\mathcal{E}(\hat{f}) \leq O(n^{-\tau_1} \log n).$$

In the case when  $\xi > 0$  and  $\alpha > 1/2$  (i.e. when the number of optimization iterations is larger than  $\sqrt{n}$  and  $R_\phi^*$  converges to  $R_{\text{hinge}}^*$  faster than  $1/\gamma$ ), we have  $\tau_1 > 1/2$ , implying that using a smooth convex loss will lead to a generalization error bound better than  $O(n^{-1/2})$  when the number of training examples is limited. This implies that for smooth loss function to achieve a binary excess error to the extent which is achievable by corresponding non-smooth loss we can run the first order optimization method for a less number of iterations. This is because our result examines the binary excess risk by taking into account the optimization complexity.

We also note  $1/(2\tau_1 - 2\tau_2)$  is given by

$$\frac{1}{2\tau_1 - 2\tau_2} = \frac{1 + \xi}{1/2 + \xi \min(1, 2\alpha - 1)}$$

When  $\alpha \leq 3/4$ , we have  $n_0 \geq K_3(B, \delta)[R_{\text{hinge}}^*]^{-2}$ , which could be a large number when  $R_{\text{hinge}}^*$  is very small.

## 6 Conclusion

In this paper we have investigated how the smoothness of loss function being used as the surrogate of 0-1 loss function in empirical risk minimization affects the excess binary risk. While the relation between convex excess risk and binary excess risk being provably established previously under weakest possible condition such as differentiability, it was not immediately obvious how smoothness of convex surrogate relates to statistical consequences. This paper made first step towards understanding this affect. In particular, in contrast to optimization and generalization analysis that favor smooth surrogate losses, our results revealed that smoothness degrades the binary excess risk. To investigate guarantees on which the smoothness would be a desirable property, we proposed a unified analysis that integrates errors in optimization, generalization, and translating convex excess risk into binary excess risk. Our result shows that under favorable conditions and with appropriate choice of smoothness parameter, a smoothed hinge loss can achieve a binary excess risk that is better than  $O(1/\sqrt{n})$ .

## Appendix A. Proof of Theorem 2

We first compute

$$z = \arg \min_{z'} \frac{1+\eta}{2} \phi(z'; \gamma) + \frac{1-\eta}{2} \phi(-z'; \gamma)$$

By setting the derivative to be zero, we have

$$\frac{1+\eta}{1+\exp(-\gamma(1-z))} = \frac{1-\eta}{1+\exp(-\gamma(1+z))}$$

and therefore

$$(1+\eta) \exp(-\gamma z) - (1-\eta) \exp(\gamma z) + 2\eta \exp(\gamma) = 0.$$

Solving the equation, we obtain

$$\exp(-\gamma z) = \frac{-\eta \exp(\gamma) + \sqrt{\eta^2 \exp(2\gamma) + (1-\eta^2)}}{1+\eta}$$

and

$$\exp(\gamma z) = \frac{\eta \exp(\gamma) + \sqrt{\eta^2 \exp(2\gamma) + (1-\eta^2)}}{1-\eta}.$$

It is easy to verify that  $\text{sgn}(z) = \text{sgn}(\eta)$ . This is because if  $\eta > 0$ , we have

$$\exp(-\gamma z) \leq \frac{\sqrt{1-\eta^2}}{1+\eta} = \sqrt{\frac{1-\eta}{1+\eta}} < 1$$

and therefore  $z > 0$ . On the other hand, when  $\eta < 0$ , we have

$$\exp(\gamma z) = \frac{1+\eta}{-\eta \exp(\gamma) + \sqrt{\eta^2 \exp(2\gamma) + (1-\eta^2)}} \leq \sqrt{\frac{1+\eta}{1-\eta}} < 1,$$

and therefore  $z < 0$ . Using the solution for  $z$ , we compute  $\phi(\eta)$  as

$$\begin{aligned} \psi(\eta; \gamma) &= \frac{1+\eta}{2} \phi(z; \gamma) + \frac{1-\eta}{2} \phi(z; \gamma) - \min_z \frac{1+\eta}{2} \phi(z; \gamma) + \frac{1-\eta}{2} \phi(z; \gamma) \\ &= -\frac{1+\eta}{2\gamma} \log \frac{1+\exp(\gamma(1-z))}{1+\exp(\gamma)} - \frac{1-\eta}{2\gamma} \log \frac{1+\exp(\gamma(1+z))}{1+\exp(\gamma)}. \end{aligned}$$

By defining constants  $C_1 = -\eta e^\gamma + \sqrt{\eta^2 e^{2\gamma} + 1 - \eta^2}$  and  $C_2 = \eta e^\gamma + \sqrt{\eta^2 e^{2\gamma} + 1 - \eta^2}$ , we can rewrite the transform function  $\psi(\eta; \gamma)$  as

$$\psi(\eta; \gamma) = -\frac{1+\eta}{2\gamma} \log \left( \frac{1}{1+e^\gamma} \left[ 1 + e^\gamma \frac{C_1}{1+\eta} \right] \right) - \frac{1-\eta}{2\gamma} \log \left( \frac{1}{1+e^\gamma} \left[ 1 + e^\gamma \frac{C_2}{1-\eta} \right] \right).$$

## Appendix B. Proof of Theorem 3

Since the expression for  $\psi(\eta; \gamma)$  is symmetric in terms  $\eta$ , we will only consider the case when  $\eta > 0$ . First, we have

$$\frac{C_1 e^\gamma}{1+\eta} = \frac{1-\eta}{\eta + \sqrt{\eta^2 + (1-\eta^2)e^{-2\gamma}}} \leq \frac{1-\eta}{2\eta}.$$

Similarly, we have

$$\frac{C_2 e^\gamma}{1-\eta} = \frac{e^\gamma}{1-\eta} \left( \eta e^\gamma + \sqrt{\eta^2 e^{2\gamma} + 1 - \eta^2} \right) \leq \frac{1+\eta}{1-\eta} e^{2\gamma}$$

Thus, we have

$$\begin{aligned}
\psi(\eta; \gamma) &\geq \frac{1+\eta}{2\gamma} \log(1+e^\gamma) - \frac{1+\eta}{2\gamma} \log\left(\frac{1-\eta}{2\eta}\right) - \frac{1-\eta}{2\gamma} \log\left(\frac{1+\eta}{1-\eta} e^\gamma\right) \\
&\geq \eta - \frac{1+\eta}{2\gamma} \log\left(\frac{1-\eta}{2\eta}\right) - \frac{1-\eta}{2\gamma} \log\left(\frac{1+\eta}{1-\eta}\right) \\
&\geq \eta - \frac{1}{\gamma} \log\left(\frac{1-\eta^2}{4\eta} + \frac{1+\eta}{2}\right) = \eta - \frac{1}{\gamma} \log\left(\frac{1}{4\eta} + \frac{\eta}{4} + \frac{1}{2}\right)
\end{aligned}$$

where the last inequality follows from the concaveness of  $\log(\cdot)$  function. As a result when  $\eta \in (-1, 1)$  we have

$$\frac{1}{4\eta} + \frac{\eta}{4} + \frac{1}{2} \leq \frac{1}{\eta},$$

which completes the proof.

## Appendix C. Proof of Theorem 5

Applying Lemmas 1 and 2 to the solution to the empirical convex risk minimizer  $\hat{f}$ , we have

$$\begin{aligned}
R_\phi(\hat{f}) &\leq \hat{R}_\phi(\hat{f}) + K_1 \left( \frac{(B + \gamma B^2)t}{n} + \sqrt{\hat{R}_\phi(\hat{f}) \frac{(B + \gamma B^2)t}{n}} \right) \\
&\leq \hat{R}_\phi(f_\gamma^*) + \frac{\gamma B^2}{(k+2)^2} + K_1 \left( \frac{(B + \gamma B^2)t}{n} + \sqrt{\hat{R}_\phi(f_\gamma^*) \frac{(B + \gamma B^2)t}{n}} + \sqrt{\frac{\gamma B^2 (B + \gamma B^2)t}{(k+2)^2 n}} \right)
\end{aligned} \tag{7}$$

On the other hand, by the application of the Bernstein's inequality [4], with probability at least  $1 - \delta$  we have

$$\begin{aligned}
\hat{R}_\phi(f_\gamma^*) - R_\phi(f_\gamma^*) &\leq \frac{4B \log \frac{1}{\delta}}{n} + \sqrt{\frac{4\mathbb{E}_{(\mathbf{x}, y) \sim P_{\mathbf{x} \times \mathbf{y}}} \left[ (\phi(y f_\gamma^*(\mathbf{x}); \gamma) - R_\phi(f_\gamma^*))^2 \right] \log \frac{1}{\delta}}{n}} \\
&\leq \frac{4B \log \frac{1}{\delta}}{n} + \sqrt{\frac{8B R_\phi(f_\gamma^*) \log \frac{1}{\delta}}{n}}.
\end{aligned} \tag{8}$$

We conclude the proof by plugging in (7) with (8), replacing the constants with a new universal constant  $K$ , and noting that  $t = \log \frac{1}{\delta} + \log^3 n$ .

## References

- [1] Sanjeev Arora, László Babai, Jacques Stern, and Z Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. In *Foundations of Computer Science, 1993. Proceedings., 34th Annual Symposium on*, pages 724–733. IEEE, 1993.
- [2] Peter L. Bartlett, Michael I. Jordan, and Jon D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [3] Shai Ben-David, David Loker, Nathan Srebro, and Karthik Sridharan. Minimizing the misclassification error rate using a surrogate convex loss. In *ICML*, 2012.
- [4] Stéphane Boucheron, Gábor Lugosi, and Olivier Bousquet. Concentration inequalities. In *Advanced Lectures on Machine Learning*, pages 208–240. Springer, 2004.
- [5] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

- [6] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *Computational learning theory*, pages 23–37. Springer, 1995.
- [7] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *The annals of statistics*, 28(2):337–407, 2000.
- [8] Klaus-Uwe Hoffgen, Hans-Ulrich Simon, and Kevin S Vanhorn. Robust trainability of single neurons. *Journal of Computer and System Sciences*, 50(1):114–125, 1995.
- [9] Wenxin Jiang. Process consistency for adaboost. *The Annals of Statistics*, 32(1):13–29, 2004.
- [10] Adam Tauman Kalai, Adam R Klivans, Yishay Mansour, and Rocco A Servedio. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6):1777–1805, 2008.
- [11] Adam Tauman Kalai and Ravi Sastry. The isotron algorithm: High-dimensional isotonic regression. In *COLT*, 2009.
- [12] Yi Lin. A note on margin-based loss functions in classification. *Statistics & probability letters*, 68(1):73–82, 2004.
- [13] Gábor Lugosi and Nicolas Vayatis. On the bayes-risk consistency of regularized boosting methods. *Annals of Statistics*, pages 30–55, 2004.
- [14] Enno Mammen and Alexandre B Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.
- [15] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer, 2004.
- [16] Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. Smoothness, low noise and fast rates. In *Advances in Neural Information Processing Systems*, pages 2199–2207, 2010.
- [17] Ingo Steinwart. Consistency of support vector machines and other regularized kernel classifiers. *Information Theory, IEEE Transactions on*, 51(1):128–142, 2005.
- [18] Paul Tseng. On accelerated proximal gradient methods for convex-concave optimization, 2009.
- [19] Vladimir N Vapnik. Statistical learning theory. 1998.
- [20] Tong Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, pages 56–85, 2004.
- [21] Ding-Xuan Zhou. Capacity of reproducing kernel spaces in learning theory. *Information Theory, IEEE Transactions on*, 49(7):1743–1752, 2003.